

CHARACTERIZATION AND COMPUTATION OF INFINITE HORIZON SPECIFICATIONS OVER MARKOV PROCESSES

Ilya Tkachev Alessandro Abate

Delft University of Technology

i.tkachev@tudelft.nl a.abate@tudelft.nl

ABSTRACT. The work is devoted to the formal verification of specifications over general discrete-time Markov processes, with an emphasis on infinite-horizon properties. These properties, formulated in a modal logic known as PCTL, can be expressed through value functions over the state space. The main goal is to understand how structural features of the model (primarily the presence of absorbing sets) influence the uniqueness of the solutions of corresponding Bellman equations. Furthermore, this contribution shows that the investigation of these structural features leads to new computational techniques to calculate the specifications of interest: the emphasis is to derive approximation techniques with associated explicit convergence rates and formal error bounds.

1. INTRODUCTION

The use of formal verification notions and methods for dynamical systems has recently become an active area of research in systems and control theory [Tab09]. One of the most efficient techniques is model-checking, which aims at determining the satisfaction set of a given specification, i.e. the set of all states that initialize realizations verifying that specification. Probabilistic Computation Tree Logic (PCTL) is a modal logic which is widely used in formal verification and dependability analysis to express specifications for discrete-time probabilistic processes [BK08, Chapter 10]. The special case of discrete-time Markov Chains (dtMC) – models over discrete (countable) spaces – is well-studied in the literature and PCTL specifications can be verified over these models in an automatic manner by employing computationally advantageous probabilistic model checking techniques [HKNP06, KKZ05]. PCTL model checking has also been validated over numerous compelling applications [FKNP11].

The formal extension of PCTL to discrete-time Markov processes (dtMP) over general (uncountable) state spaces has only recently been discussed in [Hut05, RCSL10]. The latter work in particular has expressed the satisfaction set of a given PCTL specification as the level set of an associated state-dependent value function, and has further characterized the computation of such value function via dynamic programming (DP) [BS78]. Within PCTL, there is a clear distinction between finite-horizon specifications (the satisfiability of which depends on finite realizations of the system) and infinite-horizon specifications (those characterized over infinite paths). In the context of dtMC with a finite state space, DP over a finite horizon is performed by iterative matrix multiplications, whereas DP over an infinite horizon is reduced to solving systems of linear equations. On the other hand, over a general state space the corresponding procedures – namely Bellman iterations and Bellman equations – involve integral operators. Recent work (see e.g. [APLS08]) has shown that explicit analytical solutions over uncountable state-spaces are not to be found in general, and has stressed the need for methods to compute value functions with any given precision.

In the context of dtMP, the work in [Hut05] has put forward finite abstractions [Pap03], where measures are approximated by monotone functions of sets. Although these abstractions are sound and upper and lower bounds for the expression of value functions have been derived [Hut05, Theorem 33], no method to tune them has been given. Also, their tightness and usefulness or possible triviality (i.e. conditions for the errors to lie within $(0, 1)$) has not been addressed. [RCSL10], in turn, has characterized PCTL specifications and their associated value functions with an emphasis on the issue of uniqueness of solutions of the related Bellman equations. The following questions have been left open to investigation:

- (1) how to compute finite-horizon value functions in PCTL with any given precision?
- (2) since in general value functions are not known exactly and satisfaction sets are expressed as level set of these function, how to verify nested PCTL formulae (namely, specifications where the satisfaction set for the first formula appears in the definition of a second one)?
- (3) how to verify infinite-horizon PCTL specifications in PCTL, particularly if the sufficient conditions for the uniqueness of solutions of Bellman equations in [RCSL10] are not satisfied?

With focus on 1), finite-horizon computations have recently received considerable attention. For discrete-time Stochastic Hybrid Systems (a class of dtMP), the work in [AKLP10] has put forward finite abstraction techniques to perform DP iterations over corresponding finite state-space dtMC. These results have been further sharpened in [SA11], where abstractions by state-space partitioning are obtained adaptively, in accordance to a specification-dependent error. In both works the explicit abstraction error grows linearly with the time horizon of the corresponding PCTL specification, which does not allow applying the developed methods directly to the verification of infinite-horizon properties.

This contribution of this work is hence focused on questions 2) and 3) and is twofold: the first goal is to complete the formal discussion on general state-space PCTL verification by dealing with nested formulae; the second goal (and the main task of this work) is to provide both analysis and computational tools for infinite-horizon PCTL specifications under conditions on the model that are as weak as possible and that are easy to verify.

In order to address question 2), we introduce the concepts of sub- and super-satisfaction sets for PCTL specifications, the characterization of which requires only approximate knowledge of the corresponding value functions. Specifically, we show how the sub- and super-satisfaction sets of a nested sub-formula propagate to the corresponding sets of the main formula: this is achieved by using monotonicity properties of corresponding value functions.

In order to tackle question 3), we extend and generalize recent results in [TA11, TA12], showing that the sufficient condition provided in [RCSL10] for the uniqueness of the solution of a Bellman equation is only satisfied if the solution is trivial in some sense. We further show that a weaker version of this condition is both necessary and sufficient if the dtMP admits certain continuity properties. This result leads to novel techniques to solve Bellman equations whenever their solution is not unique, and provides approximation techniques with associated explicit convergence rates and error bounds. These techniques are based on the reduction of the infinite-horizon problem to a finite-horizon one, for which computational methods available in the literature [AKLP10, SA11] can be directly applied. We furthermore discuss the relationship between the issue of uniqueness of solution and the presence of absorbing sets over the (uncountable) state space: absorbing sets are shown to play a fundamental role for both the characterization and the computation of infinite-horizon properties.

The contribution is organized as follows. Section 2 introduces discrete-time Markov processes and PCTL specifications, and discusses the verification of nested PCTL formulae. Section 3 dives in depth into infinite-horizon problems. Section 4 provides two case studies to discuss the results and finally Section 5 concludes the work.

Throughout the article we use tools of measure theory and of functional analysis. The following references can be consulted: [Dur04] for probability theory, [Rev84] for Markov processes and their kernel representation, and [Rud87] for functional analysis and measure theory.

2. MARKOV PROCESSES AND PCTL

2.1. Discrete-time Markov processes. Let $(\mathcal{X}, \mathcal{B})$ be a measurable space and let $P : \mathcal{X} \times \mathcal{B} \rightarrow [0, 1]$ be a stochastic kernel, so that $P(\cdot, B)$ is a non-negative measurable function for any set $B \in \mathcal{B}$ and $P(x, \cdot)$ is a probability measure on $(\mathcal{X}, \mathcal{B})$ for any $x \in \mathcal{X}$. The space of trajectories is denoted by $\Omega = \mathcal{X}^{\mathbb{N}_0}$ (here $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$) and its product σ -algebra with \mathcal{F} . It follows from [Rev84, Theorem 2.8] that there exists a unique discrete-time Markov process (dtMP) $X = (X_n)_{n \geq 0}$ with the transition kernel P , that is, for any $x \in \mathcal{X}$ there exists a unique probability measure P_x on (Ω, \mathcal{F}) such that $P_x(X_0 = x) = 1$, and for any measurable set $B \in \mathcal{B}$ and any time epoch $n \geq 0$

$$(2.1) \quad P_x(X_{n+1} \in B | X_0, X_1, \dots, X_n) = P(X_{n+1} \in B | X_n) = P(X_n, B).$$

Equation (2.1) characterizes the Markov property and it indicates that the future of the process X_{n+1} is independent of its past history (X_0, \dots, X_{n-1}) , given its current value X_n . As a result, any dtMP can be characterized equivalently by the triple $(\mathcal{X}, \mathcal{B}, P)$.

A familiar class of dtMP is that of stochastic dynamical systems. If $(\xi_n)_{n \geq 0}$ is a sequence of iid random variables and $f : \mathcal{X} \times \mathbb{R} \rightarrow \mathcal{X}$ is a measurable map, then

$$(2.2) \quad X_{n+1} = f(X_n, \xi_n), \quad X_0 = x \in \mathcal{X},$$

is always a Markov process characterized by a kernel $Q(x, A) = \nu(\xi \in \mathbb{R} : f(x, \xi) \in A)$, where ν is the distribution of ξ_0 . Conversely, any dtMP X admits a dynamical representation as in (2.2), for an appropriate choice of the function f [Kal02, Proposition 8.6]. However, theoretical studies of dtMP, as well as the current article, usually employ the representation via stochastic kernels. The reader interested in further discussions about modeling aspects of dtMP is referred to [Mey08, Appendix A1]. Among other models related to dtMP, Labeled Markov processes [DGJP04] may also be of interest.

2.2. Probabilistic Computation Tree Logic (PCTL). PCTL is a modal logic employed to characterize classes of temporal properties of dtMC [BK08] and of dtMP [Hut05, RCSL10]. Properties are expressed as formulae in PCTL and are constructed according to the grammar of this logic. The grammar is based on AP, the set of *atomic propositions*, which can be thought of as tags or labels associated to the states of the model. Let $A \in \text{AP}$ and $x \in \mathcal{X}$; we write $x \models A$ if the atomic proposition A is valid at state x . Since there is no substantial difference between A and its *satisfaction set* $\{x \in \mathcal{X} : x \models A\} \subseteq \mathcal{X}$, we define atomic propositions to be measurable subsets of \mathcal{X} , or equivalently $\text{AP} \subseteq \mathcal{B}$, and require that $\mathcal{X} \in \text{AP}$. The grammar of PCTL is defined as follows. Atomic propositions are basic formulae that are used to build more complex formulae via logical rules. PCTL *state formulae* are subsets of \mathcal{X} , whereas *path formulae* are subsets of Ω . More precisely:

- **true** is a formula with the whole \mathcal{X} as its satisfaction set;
- each atomic proposition $A \in \text{AP}$ is a formula with A itself as its satisfaction set;
- if A and B are formulae, then so are $\neg A$ and $A \wedge B$;
- if ϕ is a path formula and $p \in [0, 1]$, then $\mathbb{P}_{\bowtie p}[\phi]$ is a (state) formula, where \bowtie can be any symbol from the collection $\{\leq, <, \geq, >\}$;

- if A and B are formulae and $n \in \mathbb{N}_0$, then XA , $A \mathsf{U}^{\leq n} B$, and $A \mathsf{U} B$ are path formulae.

The semantics of PCTL state formulae is given as follows:

$$\begin{aligned} x \models \mathsf{true} & \quad \text{for all } x \in \mathcal{X} \\ x \models A & \quad \Leftrightarrow \quad x \in A \\ x \models \neg A & \quad \Leftrightarrow \quad x \in A^c := \mathcal{X} \setminus A \\ x \models A \wedge B & \quad \Leftrightarrow \quad x \in A \cap B \\ x \models \mathbb{P}_{\bowtie p}[\phi] & \quad \Leftrightarrow \quad P_x(\phi) \bowtie p \end{aligned}$$

With regards to path formulae, the meaning of XA (the *next* operator) is $X_1 \in A$, thus $x \models \mathbb{P}_{\bowtie p}[\mathsf{XA}]$ if and only if $P(x, A) \bowtie p$. The two additional path formulae depend on the *bounded until* operator $\mathsf{U}^{\leq n}$ and on the *unbounded until* operator U . In order to characterize them through subsets of Ω , let us introduce for any set $A \in \mathcal{B}$

$$\tau_A := \inf\{n \geq 0 : X_n \in A\}$$

to be the first hitting time of a set A over a realization X_0, X_1, \dots . Clearly, τ_A is a random variable with values in $\mathbb{N}_0 \cup \{\infty\}$. We define $A \mathsf{U}^{\leq n} B = \{\tau_B \leq \tau_{A^c}, \tau_B \leq n\} \in \mathcal{F}$, whenever $A, B \in \mathcal{B}$, which means that the path formula is satisfied over a trajectory for which B holds at least once within the n -step horizon, while A is persistently valid until that moment. Similarly, for the infinite-horizon case, we define

$$A \mathsf{U} B = \{\tau_B \leq \tau_{A^c}, \tau_B < \infty\}.$$

To characterize satisfaction sets for until operators, we introduce the so called *reach-avoid* value functions: for any $A, B \in \mathcal{B}$, let us define

$$w_n(x; A, B) := P_x(A \mathsf{U}^{\leq n} B), \quad w(x; A, B) := P_x(A \mathsf{U} B),$$

which leads to expressing $\mathbb{P}_{\bowtie p}[A \mathsf{U}^{\leq n} B] = \{x \in \mathcal{X} : w_n(x; A, B) \bowtie p\}$. Functions w_n, w are measurable, thus all PCTL formulae are well-defined measurable subsets of \mathcal{X} and all path formulae are elements of \mathcal{F} [RCSL10].¹

Let us provide a few examples: if A, B are PCTL formulae, then $\mathbb{P}_{\leq 0.05}[A \mathsf{U}_{<1}[XB]]$ is a PCTL formula. Likewise, $A \Rightarrow \mathbb{P}_{\geq 0.95}[A \mathsf{U} B]$ is a PCTL formula, since $A \Rightarrow B := \neg A \vee B$ and $A \vee B := \neg(\neg A \wedge \neg B)$. However, $\mathbb{P}_{>0}[(XA) \wedge (B \mathsf{U} C)]$ is not a PCTL formula, since the logical operation \wedge is defined for state formulae but not over path formulae. Furthermore, PCTL path formula $\Diamond^{\leq n} A := \mathsf{true} \mathsf{U}^{\leq n} A = \{\tau_A \leq n\}$ is known as a *reachability* event for a given set A and relates to a wide and important class of problems in systems and control [APLS08]. Its dual, the *invariance* (or *safety*) event $\Box^{\leq n} A = \neg(\Diamond^{\leq n} A^c) = \{\tau_{A^c} > n\}$, cannot be directly expressed in PCTL since the negation of path formulae is not allowed. On the other hand,

$$P_x(\tau_{A^c} > n) = 1 - w_n(x; \mathcal{X}, A^c),$$

thus one can define $\mathbb{P}_{\bowtie p}[\Box^{\leq n} A] = \mathbb{P}_{\bowtie' 1-p}[\mathsf{true} \mathsf{U}^{\leq n} A^c]$, where the symbol \bowtie' stands for \geq , the symbol \leq' stands for $>$, and vice versa. We denote the invariance value functions by

$$(2.3) \quad u_n(x; A) := 1 - w_n(x; \mathcal{X}, A^c), \quad u(x; A) := 1 - w(x; \mathcal{X}, A^c).$$

The results for reach-avoid and invariance given in this work can thus be directly exported to the reachability property. The latter represents also a crucial property for other types of logics, for instance linear temporal logic (LTL) [BK08, Chapter 5]. In particular, [AKM11] has argued that the verification of LTL specifications over a dtMP

¹ Although the theory in [RCSL10] has been developed for models with \mathcal{X} carrying a topological structure, all the results on measurability hold without this requirement and as such they are also valid in the present instance. This work resort to a topological structure over the state space only in Section 3.2.

expressed via specific automata, can be reduced to a reachability problem [AKM11, Theorem 4].

2.3. Nested PCTL properties. As mentioned in the introduction, it is in general not expected that the value functions w_n and w can be expressed explicitly. An alternative goal is the following [AKLP10]: given any precision level $\delta > 0$, find approximate functions \hat{w}_n and \hat{w} such that $|\hat{w}_n(x) - w_n(x)| \leq \delta$ and $|\hat{w}(x) - w(x)| \leq \delta$, for all $x \in \mathcal{X}$. Consider however the formula $\mathbb{P}_{\geq p_1} [A \cup \mathbb{P}_{\leq p_2} [B \cup C]]$: if the value function $w(x; B, C)$ can only be characterized approximately, what set should be considered to characterize $\mathbb{P}_{\leq p_2} [B \cup C]$? And how could this set be used in the parent formula? To resolve this issue we need the following fact.

Proposition 1. *Let $A \subseteq A^*$ and $B \subseteq B^*$ be elements of \mathcal{B} and let $n \in \mathbb{N}_0$. For all $x \in \mathcal{X}$:*

$$w_n(x; A, B) \leq w_n(x; A^*, B^*), \quad w(x; A, B) \leq w(x; A^*, B^*).$$

Proof. Since $\{\tau_B \leq \tau_{A^c}, \tau_B \leq n\} \subseteq \{\tau_{B^*} \leq \tau_{(A^*)^c}, \tau_{B^*} \leq n\}$ the proof immediately follows as the probability measure P_x is a monotonic function of sets for any $x \in \mathcal{X}$. \square

For a PCTL formula $A \in \mathcal{B}$, we say that A_* (A^*) is a *subsatisfaction* (*supersatisfaction*) set if $A_* \subseteq A$ ($A \subseteq A^*$). Clearly, A_* denotes a conservative set, the states of which also satisfy A , while A^* denotes a relaxed set: any state in $(A^*)^c$ does not satisfy A either.

As done above, let \hat{w}_n, \hat{w} denote some abstract δ -approximations of w_n and w , respectively. Let us show as an example, how the formula $\mathbb{P}_{\geq p_1} [A \cup \mathbb{P}_{\leq p_2} [B \cup C]]$ can be verified. Since

$$\hat{w}(x; B, C) - \delta \leq w(x; B, C) \leq \hat{w}(x; B, C) + \delta,$$

it follows that $\hat{w}(x; B, C) \leq p_2 - \delta$ implies $w(x; B, C) \leq p_2$, and that $\hat{w}(x; B, C) > p_2 + \delta$ implies $w(x; B, C) > p_2$. As a result, if we denote $D = \mathbb{P}_{\leq p_2} [B \cup C]$, then the sets

$$D_* = \{x \in \mathcal{X} : \hat{w}(x; B, C) \leq p_2 - \delta\}, \quad D^* = \{x \in \mathcal{X} : \hat{w}(x; B, C) > p_2 + \delta\}$$

represent sub- and super-satisfaction sets for D . Finally, from Proposition 1 we obtain:

$$E_* = \{x \in \mathcal{X} : \hat{w}(x; A, D_*) \geq p_1 + \delta\}, \quad E^* = \{x \in \mathcal{X} : \hat{w}(x; A, D^*) < p_1 - \delta\}$$

are sub- and super-satisfaction sets for $\mathbb{P}_{\geq p_1} [A \cup \mathbb{P}_{\leq p_2} [B \cup C]]$. The application of this procedure over formulae including the operator X is direct, since $P(x, \cdot)$ is a monotonic function of a set-valued argument for any $x \in \mathcal{X}$.

A general algorithm for the verification of nested formulae follows: given the ability to approximately compute value functions with a precision δ , find sub- and super-satisfaction sets for the sub-formulas on the lowest level (leaves) of a given formula tree, then use these sets to find sub- and super-satisfaction sets for higher-level formulae inductively, until the sub- and super-satisfaction sets for the given formula are found (at the root).

3. VERIFICATION OF INFINITE-HORIZON PCTL SPECIFICATIONS

The goal of this section is to investigate the verification of infinite-horizon PCTL specifications and to provide methods to compute associated value functions with any given precision. For this purpose Section 3.1 introduces DP techniques to characterize the corresponding value functions, points out related issues in their evaluation and provides sufficient conditions for the precise reduction of infinite-horizon problems to finite-horizon ones. In Section 3.2, the concept of absorbing set is used to show that for a class of problems the above conditions are also necessary, and that they relate to the uniqueness of the solution of Bellman equations. This result is further applied to derive methods to solve Bellman equations with non-unique solutions, both in the general case (which is done leveraging Lyapunov-like locally excessive functions – cfr.

Section 3.3), and in the special case of integral kernels (where such functions are not needed – cfr. Section 3.4). The presented techniques depend on the characterization of absorbing sets, which is discussed in Section 3.5.

3.1. Dynamic programming and Bellman equations. Let \mathbb{B} denote the space of all real-valued, bounded and measurable functions, which is a Banach space with a norm given by $\|f\| := \sup_{x \in \mathcal{X}} |f(x)|$ for $f \in \mathbb{B}$. An operator $\mathcal{J} : \mathbb{B} \rightarrow \mathbb{B}$ is called linear if $\mathcal{J}(\alpha f + \beta g) = \alpha \mathcal{J}(f) + \beta \mathcal{J}(g)$ for any $\alpha, \beta \in \mathbb{R}$ and $f, g \in \mathbb{B}$. The quantity

$$(3.1) \quad \|\mathcal{J}\| := \sup_{\|f\| \leq 1} \|\mathcal{J}f\|$$

is called the norm of the linear operator \mathcal{J} . We say that \mathcal{J} is a contraction if $\|\mathcal{J}\| \leq \rho < 1$. An important example of a linear operator associated to a dtMP is the transition operator $\mathcal{P} : \mathbb{B} \rightarrow \mathbb{B}$ given for any function $f \in \mathbb{B}$ by the following formula

$$\mathcal{P}f(x) = \int_{\mathcal{X}} f(y)P(x, dy).$$

Let us furthermore introduce an invariance operator \mathcal{J}_A , parameterized by a measurable set $A \in \mathcal{B}$, and given by $\mathcal{J}_A f(x) = 1_A(x)\mathcal{P}f(x)$. Clearly, \mathcal{J}_A is also a linear operator and $\mathcal{J}_{\mathcal{X}} = \mathcal{P}$. Moreover, \mathcal{J}_A is a monotone operator, which means that for all functions $f, g \in \mathbb{B}$ and any set $A \in \mathcal{B}$ it holds that $\mathcal{J}_A f(x) \leq \mathcal{J}_A g(x)$ for all $x \in \mathcal{X}$ whenever $f(x) \leq g(x)$ for all $x \in \mathcal{X}$. As an abbreviation, for a function $g : \mathcal{X} \rightarrow \mathbb{R}$ and a constant $\delta \in \mathbb{R}$ we use $\{g \leq \delta\} = \{x \in \mathcal{X} : g(x) \leq \delta\}$; a similar notation is used for any of the other symbols in the collection $\{<, \geq, >, =\}$.

Let us introduce a DP procedure for until-like specifications in PCTL. Let $A, B \in \mathcal{B}$ be given sets (equivalently, state formulae in PCTL). From [RCSL10, SL10] it follows:

$$(3.2) \quad \begin{cases} w_{n+1}(x; A, B) &= 1_B(x) + \mathcal{J}_{A \setminus B} w_n(x; A, B), \\ w_0(x; A, B) &= 1_B(x). \end{cases}$$

The computation in (3.2) involves iterations of the integral operator $\mathcal{J}_{A \setminus B}$. Results in [AKLP10, SA11, SA12] allow one to compute a piece-wise constant function approximation \hat{w}_n , which is such that $\|\hat{w}_n - w_n\| \leq \lambda n$, where the constant λ depends on the quality of the state space partitioning (see e.g. [SA11, Theorem 4]). Thus, in the remainder of this work we assume that finite-horizon problems can be solved approximately and with any given precision by any of the techniques given in the literature, and instead focus on the reduction of infinite-horizon problems to finite-horizon ones.

For infinite-horizon problems, it holds that $w(x; A, B) = \lim_{n \rightarrow \infty} w_n(x; A, B)$, where the limit is point-wise non-decreasing [RCSL10]. In [RCSL10, Lemma 5] the monotone convergence theorem is applied to $w_n \rightarrow w$, in order to show that the function w solves the fixpoint Bellman equation

$$(3.3) \quad w(x; A, B) = 1_B(x) + \mathcal{J}_{A \setminus B} w(x; A, B).$$

However the convergence of $w_n \rightarrow w$ is not necessarily uniform. Moreover, equation (3.3) may have multiple solutions: since it is an affine equation, if it does not have a unique solution then it admits infinitely many, spanning an affine subspace of \mathbb{B} . To further look into this issue we leverage value functions for invariance. As discussed above, the until specification can be used to express the invariance over a given set $A \in \mathcal{B}$. Using formulae (2.3) and (3.2) we obtain the following DP recursion

$$(3.4) \quad \begin{cases} u_{n+1}(x; A) &= \mathcal{J}_A u_n(x; A), \\ u_0(x; A) &= 1_A(x). \end{cases}$$

It easily follows that u_n converges point-wise non-increasingly to function u , thus

$$(3.5) \quad u(x; A) = \mathcal{J}_A u(x; A).$$

Clearly, the verification of the invariance specification inherits issues of non-uniform convergence and of non-uniqueness of the Bellman equation (3.5) from the until specification in (3.3). However, the Bellman equation for the invariance specification has the advantage of being linear and thus always admits the trivial solution $u \equiv 0$. Moreover, the analysis of the affine equation on a linear space can be reduced to the analysis of its homogeneous (linear) version: dealing with (3.5) leads to finding methods for solving (3.3) as well.

Remark 1. *There exists a least fixed-point characterization for the infinite-horizon value functions [RCSL10, Lemma 6]: $w(x; A, B)$ is the least non-negative solution of (3.3), i.e. if f is any other non-negative solution of (3.3), then $w(x; A, B) \leq f(x)$ for all $x \in \mathcal{X}$. As a result, $u(x; A)$ is the largest solution of (3.5) not exceeding 1. Although such characterization adds little to the computation of u and w , it results in the useful fact that $\|u\| = 1$ whenever u is non-trivial, namely whenever u is not identically equal to zero.*

A sufficient condition for the uniqueness of the solution of (3.5) is that $\|u_1(\cdot, A)\| < 1$ [RCSL10, Proposition 7], which in turn leads to the contractivity of the operator \mathcal{J}_A . While this condition may be easy to check, it can be restrictive: in this case (3.5) admits the unique solution $u \equiv 0$. As a result, any invariance problem with a non-trivial solution will not satisfy this sufficient condition. It follows that the weaker condition $\|u_n(\cdot, A)\| < 1$, for some $n \geq 1$, is also sufficient for the uniqueness of the solution of (3.5). Let us introduce the quantities

$$m(A) = \inf \{m \geq 0 : \|u_m(\cdot, A)\| < 1\}, \quad \rho(A) = \|u_{m(A)}(\cdot, A)\|,$$

for any $A \in \mathcal{B}$, where we set $\rho(A) := 1$ if $m(A) = \infty$. The quantity $m(A)$ is discussed in more detail for the special case of Markov Chains in Section 3.4.

Proposition 2. *Let $A \in \mathcal{B}$ and denote for simplicity $m := m(A)$ and $\rho := \rho(A)$. Then:*

- i. *if $m < \infty$, then $u(\cdot, A) \equiv 0$, and for all $n \geq 0$ it holds that $\|u_n(\cdot, A)\| \leq \rho^{\lfloor \frac{n}{m} \rfloor}$;*
- ii. *if $A, B \in \mathcal{B}$ are disjoint² and $m < \infty$, then for all $n \geq 0$*

$$(3.6) \quad 0 \leq w(x; A, B) - w_n(x; A, B) \leq \frac{m}{1 - \rho} \rho^{\lfloor \frac{n}{m} \rfloor}.$$

Proof. For part (i), we have from (3.4) that $u_n = (\mathcal{J}_A)^{n-k} u_k$, for all $0 \leq k \leq n$. Clearly, from the finiteness of m and the definition of ρ it follows that $u_m(\cdot, A) \leq 1_A(\cdot) \rho$, so

$$u_n(\cdot, A) \leq \rho \cdot (\mathcal{J}_A)^{n-m} 1_A(\cdot) = \rho u_{n-m}(\cdot, A).$$

for $n \geq m$. By induction we obtain that $\|u_n(\cdot, A)\| \leq \rho^{\lfloor \frac{n}{m} \rfloor}$, so that

$$u(\cdot, A) = \lim_{n \rightarrow \infty} u_n(\cdot, A) = 0.$$

For part (ii), we define functions $\Delta_n(x) := w_{n+1}(x; A, B) - w_n(x; A, B)$. Clearly, it holds that $\Delta_0(x) = 1_A(x)P(x, B)$ and $\Delta_{n+1}(x) = \mathcal{J}_A \Delta_n(x)$. Moreover, from the fact that $\Delta_0(x) \leq u_0(x; A)$ and the monotonicity of the operator \mathcal{J}_A , we have that $\Delta_n(x) \leq u_n(x; A)$. It further follows that

$$w(x; A, B) - w_n(x; A, B) = \sum_{i=n}^{\infty} \Delta_i(x) \leq \sum_{i=n}^{\infty} \rho^{\lfloor \frac{i}{m} \rfloor} \leq \sum_{k=\lfloor n/m \rfloor}^{\infty} m \rho^k = \frac{m}{1 - \rho} \rho^{\lfloor \frac{n}{m} \rfloor}.$$

□

²In the following, for the sake of the simplicity the set-valued arguments of the reach-avoid value functions are assumed to be disjoint. This assumption does not affect the generality of the results, since $w(x; A, B) = w(x; A \setminus B, B)$ and hence any reach-avoid problem can be always considered as a problem on disjoint sets.

As mentioned before, one goal of this section is to reduce a given infinite-horizon problem to a finite-horizon one, with the ability to tune the error incurred in this reduction. If $m(A) < \infty$, and since the right-hand side in (3.6) decreases exponentially fast with respect to n , Proposition 2 provides a method to achieve this. In the following, the condition $m(\cdot) < \infty$, for an appropriate set-valued argument, indicates that the corresponding infinite-horizon problem can be reduced (and thus solved).

It is worth mentioning that Proposition 2 elucidates the difficulty in the direct extension of the error bounds in [AKLP10, SA11] from finite- to infinite-horizon problems: the developed finite-horizon approximation techniques can be interpreted as providing a perturbation \tilde{P} of the original stochastic kernel P . Thus, they are tailored at rendering the one-step error $\|\tilde{P} - P\|$ (under the operator norm in (3.1)) as small as possible. However, in general a bound on the one-step error cannot be extended over an infinite time horizon, as the following argument shows. Let us consider the case where the solution of the invariance problem on a set A for the dtMP $(\mathcal{X}, \mathcal{B}, P)$ is non-trivial. We denote the corresponding value function as $u(x; A)$. It follows from Remark 1 that $\|u\| = 1$. Let ν be any probability measure on $(\mathcal{X}, \mathcal{B})$ such that $\nu(A^c) > 0$, and define $P^\delta(x; A) := (1 - \delta)P(x; A) + \delta\nu(A)$, for $\delta \in (0, 1)$. We have

$$\|P^\delta f - Pf\| = \left\| \delta \cdot \int_{\mathcal{X}} f(y) \nu(dy) - \delta \cdot Pf \right\| \leq \delta (\|f\| + \|Pf\|) \leq 2\delta \|f\|,$$

for any function $f \in \mathbb{B}$. Hence $\|P^\delta - P\| \leq 2\delta$, so that it can be made arbitrarily small. On the other hand, if we denote by u^δ the solution of the invariance problem on A for dtMP $(\mathcal{X}, \mathcal{B}, P^\delta)$, we obtain that $\|u_1^\delta\| \leq 1 - \delta < 1$. As a result, $u^\delta \equiv 0$ by Proposition 2, so that $\|u - u^\delta\| = 1$, regardless of how small δ is.

3.2. Absorbing and simple sets. From Proposition 2 it follows that the condition $m(A) < \infty$ in particular implies the uniqueness of the solution of the corresponding Bellman equation. It turns out that under some continuity assumptions on the kernel P this condition is also necessary. Before we proceed, we introduce the notion of absorbing set, which is crucial for further discussions.

Definition 1. A set $A \in \mathcal{B}$ is called *absorbing* if $P(x, A) = 1$, for all $x \in A$. If for $A \in \mathcal{B}$ there is an absorbing subset $A' \subseteq A$ such that $A'' \subseteq A'$ whenever $A'' \subseteq A$ is absorbing, then we say that A' is the *largest absorbing subset* of A and write $A' = \text{l.a.s.}(A)$. The set A is called *simple* if it does not have non-empty absorbing subsets, i.e. $\text{l.a.s.}(A) = \emptyset$, and *non-simple* otherwise.

Clearly, the whole state space \mathcal{X} and the empty set \emptyset are always absorbing, and if $(A_n)_{n \geq 0}$ is a sequence of non-empty absorbing sets, then their union $A = \bigcup_n A_n$ is absorbing and non-empty. However, it is by no means clear that $\text{l.a.s.}(A)$ exists for any given set A , since A may contain uncountably many absorbing subsets and their union may not be even measurable. Surprisingly, invariance value functions are useful to show that $\text{l.a.s.}(A)$ is always well-defined.

Lemma 1. Let $A \in \mathcal{B}$ and denote $A_n = \{u_n(\cdot; A) = 1\}$ for all $n \geq 0$, so that $A_0 = A$. Further, let $A_\infty = \bigcap_{n=0}^{\infty} A_n \in \mathcal{B}$, then for all $n \geq 0$ it holds that $A_{n+1} \subseteq A_n$ and

$$(3.7) \quad A_{n+1} = \{x \in A : P(x, A_n) = 1\}.$$

The set A_∞ admits the representation $A_\infty = \{u(\cdot; A) = 1\} = \text{l.a.s.}(A)$, i.e. it is the largest absorbing subset of A . In particular, if $m(A) < \infty$ then A is simple.

Proof. Let us first prove (3.7): for an arbitrary $x \in A_{n+1}$ it holds that

$$(3.8) \quad \int_A P(x, dy) \leq 1 = u_{n+1}(x; A) = 1_A(x) \int_{\mathcal{X}} u_n(y; A) P(x, dy) = \int_A u_n(y; A) P(x, dy).$$

Thus $\int_A (1 - u_n(y; A)) P(x, dy) = 0$ as it is non-positive from (3.8) and the integrand is non-negative. Due to the latter fact, $P(x, \{u_n(\cdot; A) = 1\}) = 1$ or equivalently $P(x, A_n) = 1$. Conversely, $x \in A$ be any state such that $P(x, A_n) = 1$ and let us show that $x \in A_{n+1}$. Indeed,

$$u_{n+1}(x; A) = \int_{\mathcal{X}} u_n(y; A) P(x, dy) \geq \int_{A_n} u_n(y; A) P(x, dy) = P(x, A_n) = 1,$$

thus $x \in A_{n+1}$. Since $A_{n+1} = \{x \in A : P(x, A_n) = 1\}$ and $A_1 \subseteq A_0$, then $A_2 \subseteq A_1$, and by induction $A_{n+1} \subseteq A_n$ for all $n \geq 0$. If $u(x; A) = 1$ for some $x \in A$, then $u_n(x; A) = 1$ and $x \in A_n$ for all $n \geq 0$, hence $x \in A_\infty$. If $x \in A_\infty$, then $x \in A_n$ for all $n \geq 0$, so that $u(x; A) = \lim_{n \rightarrow \infty} u_n(x; A) = 1$.

Suppose now that A is non-simple and that A' is its arbitrary absorbing subset. Clearly $u(x; A) = 1$ for all $x \in A'$, hence $A' \subseteq A_\infty$. Furthermore, if $A_\infty \neq \emptyset$, then for any $x \in A_\infty$ and $n \geq 0$ it holds that $x \in A_{n+1}$, hence $P(x, A_n) = 1$. This implies that A_∞ is absorbing since

$$P(x, A_\infty) = P\left(x, \bigcap_{n=0}^{\infty} A_n\right) = \lim_{n \rightarrow \infty} P(x, A_n) = 1,$$

which leads to conclude that A_∞ is the largest absorbing subset of A . \square

As it has been mentioned above, some continuity assumption on the kernel P are needed in order to sharpen the results. To do so, the state space needs to be endowed with a certain topological structure (see e.g. [HLL96]).

Definition 2. A state space $(\mathcal{X}, \mathcal{B})$ is called *topological* if \mathcal{X} is a Borel subset of a Polish (i.e. a metrizable, complete, and separable) space and if \mathcal{B} is a Borel σ -algebra of \mathcal{X} . A kernel P on a topological space is called *Feller* or *weakly continuous* if the function $\mathcal{P}f$ is upper semi-continuous (u.s.c.) whenever $f \in B$ is u.s.c. [HLL96, Appendix C].

A dtMP $(\mathcal{X}, \mathcal{B}, P)$ is said to be *weakly continuous* whenever $(\mathcal{X}, \mathcal{B})$ is a topological state space and P is weakly continuous.

The next theorem shows that for a weakly continuous dtMP, the infinite-horizon problem over a compact set A can be directly reduced to the finite-horizon one (in the sense that $m(A) < 1$) if and only if the set A is simple.

Theorem 1. Let $(\mathcal{X}, \mathcal{B})$ be a topological state space and A be a compact set. If P is weakly continuous then l.a.s.(A) is a compact set and the following statements are equivalent:

- (1) $m(A) < \infty$;
- (2) \mathcal{J}_A^n is a contraction on \mathbb{B} for some finite n (contractivity);
- (3) equation (3.5) has a unique solution (uniqueness);
- (4) $u(x; A) = 0$ for all $x \in \mathcal{X}$ (triviality);
- (5) the set A is simple: $A_\infty = \emptyset$ (simplicity).

Proof. 1) \Rightarrow 2) Clearly, for any function $f \in \mathbb{B}$ it follows that $\mathcal{J}_A f(x) \leq \|f\| 1_A(x)$ for all states $x \in \mathcal{X}$. Thus if $m(A) < \infty$, then $\mathcal{J}_A^{m(A)+1}$ is a contraction since

$$\|\mathcal{J}_A^{m(A)+1} f\| \leq \|f\| \cdot \|\mathcal{J}_A^{m(A)} 1_A\| = \|f\| \cdot \|u_{m(A)}(\cdot; A)\| \leq \rho(A) \|f\|.$$

2) \Rightarrow 3) If $f \in \mathbb{B}$ be a solution of (3.5), i.e. $f = \mathcal{J}_A f$. By induction we have $f = \mathcal{J}_A^n f$, which by contraction mapping theorem [Rud76] implies the uniqueness of the fixpoint f .

3) \Rightarrow 4) follows from the linearity of (3.5) and 4) \Rightarrow 5) from Lemma 1, so we only have to show that 5) \Rightarrow 1). Suppose this is not true, i.e. $m(A) = \infty$ but A is simple. It follows that $A_n \neq \emptyset$ for all $n \geq 0$. Since A is compact and \mathcal{X} is metrizable, A is closed and hence $u_0 = 1_A$ is u.s.c. Hence u_n is u.s.c. for all $n \geq 0$ by the weak continuity of P , which implies that all sets $A_n = \{u_n(\cdot; A) \geq 1\}$ are compact. Moreover, they are not empty and so their intersection A_∞ is compact and non-empty, which leads to a contradiction. \square

Remark 2. Within the main goal of reducing infinite-horizon problems over a set A to a finite-horizon ones, let us remark that numerical methods for finite-horizon problems leading to the computation of PCTL value functions with any given precision have been developed, up to our knowledge, only for compact subsets of finite-dimensional metric spaces [AKLP10, SA11] – this aligns with the assumption raised for Theorem 1. Also, there conditions required on the kernel P are stronger than the weak continuity raised above. Taking all of this into account, the assumptions in Theorem 1 are rather mild. Furthermore, some of the relations in the theorem are true under even weaker conditions: we postpone the discussion of these facts to Section 6 (Appendix).

Remark 3. It follows directly from Theorem 1 that if $m(A) < \infty$ then $\mathcal{J}_A^{m(A)+1}$ is a contraction and furthermore, $\|\mathcal{J}_A^{m(A)+1}\| \leq \rho(A)$.

Corollary 1 (From Theorem 1). Let $(\mathcal{X}, \mathcal{B})$ be a topological state space and A be a compact set. If $u_n(\cdot; A)$ are u.s.c. functions in the subspace topology of A [Rud87] for all $n \in \mathbb{N}_0$, then l.a.s.(A) is compact and statements 1) – 5) in Theorem 1 are equivalent.

Proof. According to the proof of Theorem 1, the implications 1) \Rightarrow 2) \Rightarrow 3) \Rightarrow 4) do not require weak continuity of P and thus hold in the current setting as well. Moreover, because functions u_n are u.s.c., then 5) \Rightarrow 1) follows directly. \square

3.3. A decomposition technique. Although Theorem 1 has been stated in terms of value functions for the invariance problem, its application to the issue of uniqueness of the solution of a reach-avoid problem is direct, since (3.5) is a homogeneous version of (3.3). As a result, if the dtMP $(\mathcal{X}, \mathcal{B}, P)$ is weakly continuous, sets A, B are disjoint and A is compact and simple, then $m(A) < \infty$ and the reach-avoid problem can be solved. Thus, the next goal is to study the case of a non-simple set A . For this objective the characterization given in Theorem 1 is again useful. We proceed assuming that the l.a.s. of a given set is known, and leave the discussion on the characterization of the l.a.s. of a given set and the verification of the simplicity of a set to Section 3.5.

If A is non-simple, the main issue preventing an efficient solution of the problem is the presence of an absorbing subset l.a.s.(A). This leads to the lack of contractivity of the operator \mathcal{J}_A and to the non-uniqueness of the solution of (3.3). Intuitively, if we were to remove some neighborhood $C \supset \text{l.a.s.}(A)$, then we would expect that $m(A \setminus C) < \infty$, so that a related problem can be solved on $A \setminus C$. Moreover, recall that the solution of the original problem on l.a.s.(A) is known: $w(x; A, B) = 0$ for all $x \in \text{l.a.s.}(A)$, since such states initialize trajectories that never reach the set B . The following result relates the solutions of the two problems:

Lemma 2 (Decomposition technique). Let sets $A, B \in \mathcal{B}$ be disjoint, and set $C \in \mathcal{B}, C \subseteq A$ be such that the invariance value function $u(\cdot; A \setminus C) \equiv 0$. Then $w(x; A \setminus C, B)$ is the unique solution of the following Bellman equation

$$(3.9) \quad w(x; A \setminus C, B) = 1_B(x) + \mathcal{J}_{A \setminus C} w(x; A \setminus C, B),$$

and for all $x \in \mathcal{X}$ the following holds:

$$(3.10) \quad 0 \leq w(x; A, B) - w(x; A \setminus C, B) \leq \sup_{y \in C} w(y; A, B).$$

Proof. As an abbreviation, let us denote $\tau_1 = \tau_{B \cup C}$ and $\tau_2 = \tau_{(A \cup B)^c}$ and let us partition the event space Ω by the following four disjoint hypotheses:

$$H_1 = \{\tau_2 < \tau_1, \tau_2 < \infty\}, \quad H_2 = \{\tau_1 = \infty, \tau_2 = \infty\},$$

$$H_3 = \{\tau_1 < \tau_2, \tau_1 < \infty, \tau_1 = \tau_B\}, \quad H_4 = \{\tau_1 < \tau_2, \tau_1 < \infty, \tau_1 = \tau_C\}.$$

Recall that $w(x; A, B) = P_x\{\tau_B < \tau_2, \tau_B < \infty\}$, thus

$$w(x; A, B) = \sum_{i=1}^4 P_x(\{\tau_B < \tau_2, \tau_B < \infty\} \cap H_i).$$

Note that the first term is zero since clearly $\{\tau_B < \tau_2, \tau_B < \infty\} \cap H_1 = \emptyset$. The second term vanishes because $P_x(H_2) = u(x; A \setminus C) \equiv 0$. Since $H_3 \subseteq \{\tau_B < \tau_2, \tau_B < \infty\}$ the third term equals to $P_x(H_3) = w(x; A \setminus C, B)$, which leaves only the fourth term to be studied. Let x be any such that $P_x(H_4) \neq 0$ and define a measure ν_x on $(\mathcal{X}, \mathcal{B})$ by

$$\nu_x(D) = P_x(X_{\tau_C} \in D | H_4), \quad D \in \mathcal{B},$$

so that clearly $\nu_x(C^c) = 0$. For such fixed x it holds that

$$\begin{aligned} 0 \leq P_x(\{\tau_B < \tau_2, \tau_B < \infty\} \cap H_4) &= P_x(\{\tau_B < \tau_2, \tau_B < \infty\} | H_4) P_x(H_4) \\ &= w(x; A, C) \cdot \int_C w(y; A, B) \nu_x(dy) \leq \sup_{y \in C} w(y; A, B). \end{aligned}$$

The same bounds clearly hold in the alternative case $P_x(H_4) = 0$.

Finally, it follows that $w(x; A \setminus C, B)$ is the unique solution of the corresponding Bellman equation (3.9) from $u(\cdot; A \setminus C) \equiv 0$ (see Proposition 5 in Section 6). \square

Corollary 2. [From Lemma 2] Let $(\mathcal{X}, \mathcal{B}, P)$ be a weakly continuous dtMP and let $A, B \in \mathcal{B}$ be disjoint and such that A is a compact, non-simple set. Let $C \subseteq A$ be an open neighborhood of l.a.s.(A) in the subspace topology of A . Then (3.10) holds for all $x \in \mathcal{X}$, and $m(A \setminus C) < \infty$.

Proof. Since C is open in A , the set $A \setminus C$ is a closed subset of a compact set A and thus itself compact. From the inclusions $\text{l.a.s.}(A) \subset C \subset A$ it follows that $A \setminus C$ is simple, hence Theorem 1 ensures that all the conditions of Lemma 2 are satisfied. \square

In order to render the result in Corollary 2 useful for the computation of the infinite-horizon reach-avoid value function, we should provide a method to choose an open neighborhood C of l.a.s.(A), such that $\sup_{y \in C} w(y; A, B) < \varepsilon$, where $\varepsilon > 0$ is a given precision level. We use the theory of excessive functions [SRG08] to achieve this goal.

Definition 3. Given a function $g \in \mathbb{B}$, the excessive set of g is $\mathcal{E}_g = \{\mathcal{P}g - g \leq 0\}$. If $\mathcal{E}_g = \mathcal{X}$, i.e. if $\mathcal{P}g(x) \leq g(x)$ for all $x \in \mathcal{X}$, then the function g is called excessive.

The relation between excessive functions and infinite-horizon invariance is given via Doob's inequality [SRG08]: if $g \in \mathbb{B}$ is an excessive, non-negative function, then

$$(3.11) \quad P_x \left\{ \sup_{n \geq 0} g(X_n) \geq \delta \right\} \leq \frac{g(x)}{\delta}.$$

for all $\delta > 0$. The inequality (3.11) can be rewritten via the invariance value function:

$$(3.12) \quad u(x; \{g < \delta\}) \geq 1 - \frac{g(x)}{\delta}.$$

Excessive functions for stochastic systems are akin to Lyapunov functions for deterministic systems, since they are characterized by decreasing behavior along the dynamics of the process, as the inequality $\mathcal{P}g \leq g$ suggests³. As is the case with Lyapunov functions for deterministic systems, it is non trivial to find excessive functions. However, it is possible to relax the assumption on global excessivity and to employ a local version of Doob's inequality.

Lemma 3. [Kus67, Theorem 12] *Let $g \in \mathbb{B}$ be a non-negative function such that for some $\delta > 0$ it holds that $\{g < \delta\} \subseteq \mathcal{E}_g$. Whenever $x \in \{g < \delta\}$, it follows that*

$$P_x \left\{ \sup_{n \geq 0} g(X_n) \geq \delta \right\} \leq \frac{g(x)}{\delta}.$$

The idea behind the proof of this lemma is to consider a set $A = \{g < \delta\}$. The related invariance value function does not depend on $P(x, \cdot)$ for $x \in A^c$, where it is simply equal to zero (recall that all the integrals in the DP recursion (3.4) are equivalently taken over the set A). As a result, exclusively the dynamics within the set A are important for the process.

Definition 4. *For a topological state space $(\mathcal{X}, \mathcal{B})$ we say that a non-negative continuous function $g \in \mathbb{B}$ is δ -locally excessive on the set $A \in \mathcal{B}$ if for some real number $\delta > 0$ it holds that $\{g = 0\} = \text{l.a.s.}(A)$ and that $\{g < \delta\} \subseteq A \subseteq \mathcal{E}_g$.*

Theorem 2. *Let $(\mathcal{X}, \mathcal{B}, P)$ be a weakly continuous dtMP and let $A, B \in \mathcal{B}$ be disjoint and such that A is a compact, non-simple set. If there exists a δ -locally excessive function g on A , then for any $\varepsilon \in (0, 1)$ it holds that $m(A \setminus \{g < \varepsilon\delta\}) < \infty$ and that*

$$(3.13) \quad 0 \leq w(x; A, B) - w(x; A \setminus \{g < \varepsilon\delta\}, B) \leq \varepsilon.$$

Proof. First, we show that for any $\varepsilon \in (0, 1)$, if $g(x) < \varepsilon\delta$ then $w(x; A, B) \leq \varepsilon$. Indeed, as $\{g < \delta\} \subseteq \mathcal{E}_g$, by Lemma 3 we have that $u(x; \{g < \delta\}) \geq 1 - \frac{g(x)}{\delta}$ for all $x \in \mathcal{X}$, so

$$u(x; A) \geq 1 - \frac{g(x)}{\delta}$$

for all $x \in \mathcal{X}$, which follows from $\{g < \delta\} \subseteq A$. Since $u(x; A) = 1 - w(x; \mathcal{X}, A^c)$, then

$$w(x; \mathcal{X}, A^c) \leq \frac{g(x)}{\delta}$$

for all $x \in \mathcal{X}$, and since $A \subseteq \mathcal{X}$ and $B \subseteq A^c$, from Proposition 1 it follows that $w(x; A, B) \leq \frac{g(x)}{\delta}$. As a result, for any $x \in \{g < \varepsilon\delta\}$ it holds that $w(x; A, B) \leq \varepsilon$.

Second, let us fix any $\varepsilon \in (0, 1)$ and denote $C = \{g < \varepsilon\delta\}$. Clearly, $\text{l.a.s.}(A) \subseteq C$ and the set $\{g < \varepsilon\delta\}$ is open in A since g is continuous on A . The statement of the theorem then follows from Corollary 2. \square

3.4. Integral kernels and discrete-space Markov Chains. From Theorem 2 it follows that for weakly continuous dtMPs a reach-avoid problem on a non-simple set can be solved if an appropriate locally excessive function is found. For a known and studied subclass of these processes the problem can be solved even without resorting to such functions. We write $P(x, dy) = p(x, y)\mu(dy)$ if P is an integral kernel with a basis μ

³From the definition of \mathcal{P} it follows that $\mathcal{P}g(x) = E_x[g(X_1)]$, where E_x denotes the expectation with respect to P_x . Thus the condition $\mathcal{P}g \leq g$ means that the expected value of the function g at the next time step is bounded by its current value, so that the function g does not increase on average along realizations of the dtMP. Thus the function $g \in \mathbb{B}$ is excessive if and only if the process $(g(X_n))_{n \geq 0}$ is a P_x -supermartingale, for all $x \in \mathcal{X}$ [PS06, p.20].

and a density p , namely when μ is a σ -finite non-negative measure, the function p is $\mathcal{B} \otimes \mathcal{B}$ -measurable⁴, and for any $A \in \mathcal{B}$ it holds that

$$P(x, A) = \int_A p(x, y) \mu(dy).$$

Furthermore, we raise the following assumption, which generalizes one used for related studies over the finite horizon [AKLP10, SA11].

Assumption 1. Let $(\mathcal{X}, \mathcal{B})$ be a topological state space and let d be a complete metric on \mathcal{X} consistent with the given topology. Assume that $\mu(A) < \infty$ for any compact set A and that there exists a \mathcal{B} -measurable function λ_A , such that $\int_A \lambda_A(y) \mu(dy) < \infty$, and such that for all $x', x'' \in A$ and μ -a.a. $y \in A$

$$(3.14) \quad |p(x'', y) - p(x', y)| \leq \lambda_A(y) \cdot d(x', x'').$$

Let us mention that if $P(x, dy) = p(x, y) \mu(dy)$, set C is such that $\mu(C) = 0$, and set A is an absorbing set, then clearly $A \setminus C$ is absorbing as well. This, in turn, implies that if A is a simple set then $A \cup C$ is also simple whenever $\mu(C) = 0$.

For a topological space \mathcal{X} , let us use the following notation: let ∂B and B° be the boundary and the interior of set $B \subseteq \mathcal{X}$.

Theorem 3. Let $(\mathcal{X}, \mathcal{B})$ be a topological space and $P(x, dy) = p(x, y) \mu(dy)$. Let sets $A, B \in \mathcal{B}$ be disjoint and let the set A be compact. Suppose that at least one of the following two items holds:

- i. P is weakly continuous;
- ii. Assumption 1.

Whenever $\mu(\partial(\text{l.a.s.}(A))) = 0$, it holds that $m(A \setminus (\text{l.a.s.}(A))^\circ) < \infty$ and for all $x \in \mathcal{X}$

$$(3.15) \quad w(x; A, B) = w(x; A \setminus (\text{l.a.s.}(A))^\circ, B).$$

Proof. To prove (3.15) we apply the decomposition technique in Lemma 2 over a set $C := (\text{l.a.s.}(A))^\circ$. The result follows since $w(y; A, B) = 0$ for all $y \in (\text{l.a.s.}(A))^\circ$, so we are only left to prove that for the set $\tilde{A} := A \setminus (\text{l.a.s.}(A))^\circ$ it holds that $m(\tilde{A}) < \infty$.

If i. is true, then $\text{l.a.s.}(A)$ is compact by Theorem 1 so

$$\text{l.a.s.}(A) = (\text{l.a.s.}(A))^\circ \cup \partial(\text{l.a.s.}(A)).$$

As a result, the set $\tilde{A} := A \setminus (\text{l.a.s.}(A))^\circ$ is compact and simple since $\mu(\partial(\text{l.a.s.}(A))) = 0$. From Theorem 1 it follows that $m(\tilde{A}) < \infty$.

If instead ii. is true, then $u_n(\cdot; A)$ is a continuous function on A for any $n \geq 0$, since

$$\begin{aligned} |u_{n+1}(x''; A) - u_{n+1}(x'; A)| &= \left| \int_A u_n(y; A) (p(x'', y) - p(x', y)) \mu(dy) \right| \\ &\leq \int_A |p(x'', y) - p(x', y)| \mu(dy) \\ &\leq \left(\int_A \lambda_A(y) \mu(dy) \right) \cdot d(x', x''), \end{aligned}$$

⁴Here $\mathcal{B} \otimes \mathcal{B}$ is the σ -algebra on the set $\mathcal{X} \times \mathcal{X}$ generated by the class of the measurable rectangles

$\mathcal{B} \times \mathcal{B} = \{B_1 \times B_2 : B_1, B_2 \in \mathcal{B}\}.$

for any pair of states $x', x'' \in A$. From Corollary 1 it follows that $\text{l.a.s.}(A)$ is compact. Following the same argument as in case i ., we obtain that the set \tilde{A} is compact and simple. Similar to (3.16) we obtain that $u_n(\cdot; \tilde{A})$ is a continuous function on \tilde{A} for any $n \geq 0$, hence $m(\tilde{A}) < \infty$ by Corollary 1. \square

Remark 4. In the special case of the invariance problem over the set A , under the assumptions of Theorem 3 it follows that $u(x; A) = w(x; A, (\text{l.a.s.}(A))^\circ)$. Moreover, the proof of Lemma 2 implies that for any initial state $x \in \mathcal{X}$, P_x -a.s. a trajectory $(X_n)_{n \geq 0}$ of the dtMP that stays invariant in the set A necessarily reaches its largest absorbing subset. Altogether, this enlightens yet another interesting relation between invariance and reach-avoid problems.

In case $\mu(\partial(\text{l.a.s.}(A))) > 0$, Theorem 3 cannot be applied and Theorem 2 is left as an alternative to tackle this problem by coming up with an appropriate δ -locally excessive function for the given set A . In general such a function may not exist even if P is a weakly continuous dtMP and A is compact (see example in Section 4.1). However, if Assumption 1 holds it follows that the function $u(\cdot; A)$ is continuous on A – the proof follows the same lines as in (3.16) – and $g(x) = 1 - u(x; A)$ is a 1-locally excessive function on A . This consideration suggests that at least under Assumption 1 there always exists a δ -locally excessive function over a compact set A .

Let us now tailor the above results to the case where the state space is countable, i.e. when the process is a discrete-time Markov Chain (dtMC). The methods developed in this section are directly applicable to the dtMC framework where, to the best of our knowledge, available techniques allow one to compute infinite-horizon reach-avoid value functions as a limit of converging iterations, but without bounds on the error [BK08, Theorem 10.15].

Without loss of generality, let us assume that $\mathcal{X} = \mathbb{N}$ is the state space of a dtMC. We endow \mathcal{X} with the discrete metric $d(i, j) := 1_{\{j\}}(i)$, so that $\mathcal{B} = 2^{\mathcal{X}}$. The basis σ -finite measure is chosen to be the counting one: $\mu(i) = 1$, for any $i \in \mathcal{X}$. Any stochastic kernel P over $(\mathcal{X}, \mathcal{B})$ can be expressed as a matrix $P = (p_{ij})_{i, j \in \mathbb{N}}$, where $p_{ij} := P(i, \{j\})$. With the chosen counting measure, the entries of the stochastic matrix P determine the density function, namely $p(i, j) = p_{ij}$. Clearly, compact subsets of \mathcal{X} correspond to finite sets, so they are of finite measure μ . Moreover, P is always weakly continuous, since on a discrete topological space every function is continuous.

Remark 5. For a dtMC the largest absorbing subset of any finite set can be found algorithmically. Indeed, from Lemma 1 it follows that $\text{l.a.s.}(A) = \{u(\cdot; A) = 1\}$, so the set can be equivalently expressed via a CTL formula: $\text{l.a.s.}(A) = \{x \in A : x \models \forall \square A\}$. As such, it can be computed in $\mathcal{O}(\mu^2(A))$ time over the dtMC graph (to be defined below) [BK08, Theorem 6.30].

Corollary 3 (from Theorem 3). Let $A, B \in \mathcal{B}$ be disjoint and A be finite, denote $\tilde{A} := A \setminus \text{l.a.s.}(A)$ and $b_i := P(i, B) = \sum_{j \in B} p_{ij}$. The reach-avoid value function $w(i; A, B)$ is defined uniquely by

$$(3.16) \quad \begin{cases} w(i; A, B) = 1 & \text{if } i \in B, \\ w(i; A, B) = b_i + \sum_{j \in \tilde{A}} p_{ij} w(j; A, B) & \text{if } i \in \tilde{A}, \\ w(i; A, B) = 0 & \text{otherwise.} \end{cases}$$

Proof. Theorem 3 holds since P is weakly continuous and the discrete topology implies that $\partial(\text{l.a.s.}(A)) = \emptyset$. Thus (3.15) holds true and the corresponding Bellman equation has the form $w(i; A, B) = 1_B(i) + 1_{\tilde{A}}(i) \sum_{j \in \mathbb{N}} p_{ij} w(j; A, B)$, which is equivalent to (3.16). \square

Note, that to find a solution for (3.16), one should solve a system of linear equations with a non-zero determinant. Moreover, notice that the square submatrix $\tilde{P} := (p_{ij})_{i,j \in \tilde{A}}$ in (3.16) is contractive since $m(\tilde{A}) < \infty$, so even for large-scale problems efficient numerical methods can be applied to solve the problem with any given precision.

Let us mention what the condition $m(A) < \infty$ means in graph-theoretical terms for a dtMC. The adjacency graph of a dtMC is a directed graph (V, E) , where with $V = \mathcal{X}$ and the set of edges E is such that $(i, j) \in E$ if and only if $p_{ij} > 0$. To an arbitrary element $i \in A$ we can assign a positive number m_i , which is the length of the shortest path in the graph (V, E) from i to A^c . Clearly, it holds that $m(A) = \sup_{i \in A} m_i$. Moreover, from this characterization it can be easily seen that $m(A) \leq \mu(A)$ if $m(A)$ is finite. As a result, if $\|u_{\mu(A)+1}(\cdot; A)\| = 1$ it follows that $m(A) = \infty$ and that

$$\text{l.a.s.}(A) = A_{\mu(A)+1} = \{u_{\mu(A)+1}(\cdot; A) = 1\}.$$

3.5. Verification of simplicity of A and determination of $\text{l.a.s.}(A)$. Let us summarize the methods developed for the solution of the infinite-horizon reach-avoid problem in the previous sections. Assume that $A, B \in \mathcal{B}$ are disjoint and let us focus on the case when the set A is compact and the kernel P is weakly continuous. If A is simple, it follows from Theorem 1 that the solution of (3.3) is unique and that $m(A) < \infty$ – the solution can be found as in Proposition 2. If A is non-simple, the solution of (3.3) is not unique and $m(A) = \infty$, thus Proposition 2 cannot be applied directly. In the latter case, there are two approaches to solve an infinite-horizon reach-avoid problem over a non-simple set A : if $P(x, dy) = p(x, y)\mu(dy)$ is an integral kernel and $\mu(\partial(\text{l.a.s.}(A))) = 0$, then Theorem 3 allows formulating an equivalent problem over the simple compact set $A \setminus (\text{l.a.s.}(A))^\circ$. Otherwise, one has to synthesize an appropriate δ -locally excessive function to apply Theorem 2.

All the instances discussed above depend on the fundamental issue of whether a given compact set A is simple or not. In general it is hard to provide an analytical answer to such a question, and no known general automatic procedure enables computing absorbing sets exactly. On the other hand, the “if and only if” nature of the results in Theorem 1 implies that this issue is not a limitation that is specific to the techniques presented in this paper: on the contrary, any other method aiming to solve a general infinite-horizon reach-avoid problem is bound to check the simplicity of a given set A .

Let us discuss instances of dtMP for which the $\text{l.a.s.}(A)$ of a given set A can be found explicitly. The case of dtMC, as discussed in Remark 5, has been recently extended to a subclass of dtMP with integral kernels in [TA12, Chapter 4.2]. In both instances all the conditions in Theorem 3 are satisfied, thus the reach-avoid problem can be solved.

Given additional knowledge on the structure of a dtMP, it may be easier to verify the dual problem, namely the simplicity of a given set A : if P is φ -irreducible [MT93, Chapter 4], then A is simple whenever $\varphi(A^c) > 0$. If φ is the maximal irreducibility measure, then A is simple if and only if $\varphi(A^c) > 0$. However, notice that for a given dtMP the verification of its irreducibility can represent an even harder requirement than the verification of the simplicity of a specific set A . Moreover, observe that any dtMP admitting two disjoint non-empty absorbing sets is not irreducible, which points out the conservatism of this condition.

An additional example where further knowledge on the structure of dtMP may shed light on the absorbance of its sets is provided in [AKM11]. As already mentioned, an automaton specification \mathcal{A} over a dtMP $\mathcal{H} = (\mathcal{X}, \mathcal{B}, P)$ can be verified as a reachability specification over the product $\mathcal{A} \times \mathcal{H}$, which is again a dtMP. The discrete structure of the automaton \mathcal{A} can be exploited in order to determine absorbing sets within the product dtMP $\mathcal{A} \times \mathcal{H}$.

Furthermore, analytical methods can be applied to find absorbing sets. If the dynamical system representation of a dtMP (2.2) is known one can try to characterize its absorbing sets, as the examples of Section 4 will display. Also, for integral kernel $P(x, dy) = p(x, y)\mu(dy)$ with density p given explicitly, one may try to check for simplicity using the following result.

Proposition 3. [TA12, Proposition 3] *For $x \in \mathcal{X}$ define $s(x) = \{y \in \mathcal{X} : p(x, y) > 0\}$. A set $A \in \mathcal{B}$ is absorbing if and only if $\mu(s(x) \setminus A) = 0$, for all $x \in A$.*

Finally, although in general the verification of the simplicity of a given set is not a decidable procedure, the following method can be applied. Let us consider the sequence $(A_n)_{n \geq 0}$ defined in Lemma 1. If $A_n = \emptyset$ for some $n \in \mathbb{N}$, then clearly A is simple. Although the definition itself requires a precise characterization of $u_n(\cdot; A)$, only the computation of $P(x, \cdot)$ is needed in (3.7), instead of consecutive integral iterations over value functions. Let us now introduce an approximate approach for the computation, using the concepts in Section 2.3: leveraging the procedure in (3.7), we have that $A_0 = A$ and that $A_{n+1} = \mathbb{P}_{\geq 1}[XA_n]$. Let us select a precision level $\delta \in (0, 1)$ and construct a sequence of supersatisfaction sets as follows:

$$A_{n+1}^* = \mathbb{P}_{\geq 1-\delta}[XA_n^*], \quad A_0^* = A.$$

By construction, $A_n \subseteq A_n^*$ for all $n \geq 0$, thus A is simple whenever $A_n^* = \emptyset$ for some $n \in \mathbb{N}$. Notice that the conditions required to implement the procedure are very general. Let us discuss its strong and weak points:

- If the exact form of P is given, then the sets A_n can be characterized explicitly. The simplicity of A is verified if the sequence $(A_n)_{n \geq 0}$ eventually contains only empty sets. On the other hand, if A is non-simple, then the set $\text{l.a.s.}(A)$ can be found whenever $A_n = A_{n+1} \neq \emptyset$ for some $n \in \mathbb{N}$. Clearly, in such a situation it holds that $\text{l.a.s.}(A) = A_n$. Finally, if A is non-simple, whenever A_{n+1} is a strict subset of A_n , one can compute $\text{l.a.s.}(A) = A_\infty$ as an intersection of the sets A_n .
- If only an approximate characterization of P is available, the simplicity of set A can be verified for sufficiently small δ and sufficiently large n . However, it is not clear how big n should be taken to ensure that $A_n^* = \emptyset$ for a given precision level δ . Due to this reason, it is extremely important to have an *a priori* upper-bound on $m(A)$, provided the latter is finite (cfr. the discussion on $m(A)$ for dtMC in Section 3.4). Furthermore, the non-simplicity of set A cannot be verified: because of the errors in the computation of A_n^* , the case $A_n = A_{n+1} \neq \emptyset$ cannot be exactly characterized.

Let us further remark that, for all $n \geq 0$, A_n^* provide an overestimation for the set $\text{l.a.s.}(A)$. If a non-trivial underestimation is available as well, the following result can be established.

Proposition 4. *Under Assumption 1, let A be a compact, non-simple set and let sets $C, D \in \mathcal{B}$ be such that $C \subseteq \text{l.a.s.}(A) \subseteq D$, where D is open. Define the sequence \tilde{u}_n as follows:*

$$(3.17) \quad \begin{cases} \tilde{u}_{n+1}(x) &= \tilde{u}_0(x) + 1_A(x) \cdot \mathcal{P}[1_{A \setminus D}(x)\tilde{u}_n(x)], \\ \tilde{u}_0(x) &= 1_A(x)P(x, D). \end{cases}$$

For any $n \geq 0$ and $x \in \mathcal{X}$ it holds that

$$(3.18) \quad \|u(\cdot; A) - \tilde{u}_n(\cdot)\| \leq \frac{m+2}{1-\rho} \cdot \rho^{\lfloor \frac{n}{m+2} \rfloor} + \frac{\alpha}{1-\rho} \cdot \frac{\alpha^{m+1} - 1}{\alpha - 1} \cdot \mu(D \setminus C).$$

where $m := m(A \setminus D) < \infty$, $\rho := \rho(A \setminus D) < 1$, and $\alpha := \sup \{p(x, y) | x, y \in A\}$.

Proof. The proof is given in [TA12, Section 3.4]. Shortly, the functional sequence in (3.17) is designed to approximate to solution of the invariance problem. The first term in the right-hand side of (3.18) comes from the upper-bound on the difference $\|u(\cdot; A) - \tilde{u}(\cdot)\|$ where $\tilde{u} := \lim_{n \rightarrow \infty} \tilde{u}_n$ (see [TA12, Theorem 3]) and the second term comes from the upper-bound on the difference $\|\tilde{u} - \tilde{u}_n\|$ (see [TA12, Proposition 2]). \square

Proposition 4 can be easily extended to be valid over reach-avoid value functions as well as dtMP with arbitrary integral kernels, so that Assumption 1 in the statement can be relaxed. However, this goal is not pursued in this paper since no procedure to find an under-approximating set C is available up to our knowledge. Indeed, any method giving a non-empty candidate for C would directly establish the non-simplicity of A .

We conclude the discussion in this section with the following practical observation: in practice stochastic kernels for a dtMP either are extracted from finite data coming from measurement experiments, or derived from some underlying analytical model. In the latter case, the model gives an additional knowledge on the structure of a dtMP which can be further used along the lines discussed in this section to find the largest absorbing subset of a given set or to verify the simplicity of such set. Conversely, when no underlying model is known and kernels are interpolated exclusively from measurements data, any kernel resulted via an interpolation technique can be negligibly perturbed in order to yield absence of absorbing subsets of given compact sets (see discussion after Proposition 2).

4. CASE STUDIES

4.1. A one-dimensional affine Gaussian system. Let $\mathcal{X} = \mathbb{R}$ be endowed with the standard topology and let \mathcal{B} be its Borel σ -algebra. Consider a sequence $(\xi_n)_{n \geq 0}$ of iid standard normal random variables and define a dtMP as

$$(4.1) \quad X_{n+1} = (\alpha + \mu X_n) + (\beta + \sigma X_n) \cdot \xi_n,$$

where $\alpha, \beta, \mu, \sigma \in \mathbb{R}$ are parameters and $X_0 = x \in \mathbb{R}$. In order to study the probabilistic invariance problem for this affine Gaussian model, let us select a compact set A in \mathbb{R} . Let us focus on how the structure of the dynamics are affected by the choice of the parameters. In order to avoid trivial constant dynamics, let us assume that at least one of the parameters $\alpha, \beta, \mu - 1, \sigma$ is non-zero. If $\beta + \sigma X_n \neq 0$ the distribution of X_{n+1} admits the whole state space \mathbb{R} as its support, thus for A to be non-simple it is necessary that point $\kappa := -\frac{\beta}{\sigma} \in A$. We then assume that $\sigma \neq 0$, since clearly if $\sigma = 0$ any compact set is simple. Moreover, for A to be non-simple the state κ has to be absorbing, so from (4.1) it must hold that $\kappa = \alpha + \mu\kappa$, so $\alpha = (1 - \mu)\kappa$. Since by Theorem 1 the solution of the invariance problem on simple sets is trivial, we focus on the case when κ is absorbing and select the parameters $\alpha := (1 - \mu)\kappa$, $\beta := -\sigma\kappa$, where $\kappa \in \mathbb{R}$ is an arbitrary state. The update equation (4.1) takes the new form:

$$X_{n+1} - \kappa = \mu(X_n - \kappa) + \sigma(X_n - \kappa)\xi_n,$$

and by applying a shift on κ , without loss of generality we can focus on the following model:

$$(4.2) \quad X_{n+1} = \mu X_n + \sigma X_n \cdot \xi_n.$$

In the latter equation σ can be assumed to be positive, since ξ_n has a symmetric distribution. The kernel associated to the dtMP (4.2) is weakly continuous and takes the following form:

$$P(x, A) = \begin{cases} \frac{1}{\sigma|x|\sqrt{2\pi}} \int_A e^{-\frac{t^2}{2(\sigma x)^2}} dt & , \text{ if } x \neq 0, \\ 1_A(0) & , \text{ if } x = 0. \end{cases}$$

Since the compact set A is non-simple if and only if $0 \in A$, let us consider the invariance problem for the set $A = [-1, 1]$. The discussion above suggests that $\text{l.a.s.}(A) = \{0\}$, so $u(0; A) = 1$. For $x \neq 0$, let us relate the original process X to the random walk (see e.g. [Dur04, Chapter 4].) Define $Y_n := \log |X_n|$, so the update equation becomes:

$$Y_{n+1} = Y_n + \log |\mu + \sigma \xi_n|,$$

where $Y_0 = y := \log |x|$. The expected value of the increment of the random walk

$$h(\mu, \sigma) := \mathbb{E} \log |\mu + \sigma \xi_1|$$

determines its asymptotic behavior. In particular, $\limsup_{n \rightarrow \infty} Y_n = +\infty$ a.s. if $h(\mu, \sigma) \geq 0$, and $\lim_{n \rightarrow \infty} Y_n = -\infty$ a.s., if $h(\mu, \sigma) < 0$ [Dur04, Chapter 4]. As a result, if the values of the parameters μ, σ are such that $h(\mu, \sigma) \geq 0$, we obtain that, for any $x \neq 0$, the following holds:

$$u(x; A) = \mathbb{P}_x \left\{ \sup_{n \geq 0} |X_n| \leq 1 \right\} = \mathbb{P}_y \left\{ \sup_{n \geq 0} \log |X_n| \leq 0 \right\} = 0,$$

which allows to conclude that so that in this case $u(x; A) = 1_{\{0\}}(x)$. We are left with the case $h(\mu, \sigma) < 0$. If we represent the kernel in the integral form $P(x, dy) = p(x, y)\mu(dy)$, we obtain $1 = P(0, \{0\}) = p(0, 0)\mu(\{0\})$, so necessarily

$$\mu(\{0\}) = \mu(\partial(\text{l.a.s.}(A))) > 0,$$

thus Theorem 3 cannot be applied. We then resort to Theorem 2, which requires finding a δ -locally excessive function.

Let us fix μ, σ and consider $g_q(x) := |x|^q$, for $q \geq 0$. If we define $b(q) = \mathbb{E}|\mu + \sigma \xi_1|^q$, then clearly $\mathcal{P}g_q(x) = b(q) \cdot g_q(x)$, so g_q is δ -locally excessive if and only if $b(q) < 1$. We obtain $b(0) = 1$ and $b'(0) = h(\mu, \sigma)$. Recall that we are now interested in the case $h(\mu, \sigma) < 0$, which leads to conclude that there always exists a $q > 0$ such that the function $g_q(x) = |x|^q$ is δ -locally excessive. Hence, for $h(\mu, \sigma) < 0$ and such q , Theorem 2 can be applied to find the solution of the invariance problem using g_q as a 1-locally excessive function. More precisely, according to (3.13) adapted to the special case of the invariance problem, we obtain:

$$0 \leq w(x; A, (-\sqrt[q]{\varepsilon}, \sqrt[q]{\varepsilon})) - u(x; A) \leq \varepsilon,$$

and function $w(x; A, (-\sqrt[q]{\varepsilon}, \sqrt[q]{\varepsilon}))$ can be computed, since $m(A \setminus (-\sqrt[q]{\varepsilon}, \sqrt[q]{\varepsilon})) < \infty$ as it follows from Theorem 2.

Finally, let us add a comment on the lack existence of a δ -locally excessive function for a weakly continuous dtMP. Consider the case in (4.2) with parameters $h(\mu, \sigma) \geq 0$, so that $u(x; A) = 1_{\{0\}}(x)$. If there existed a function g that is δ -locally excessive on $A = [-1, 1]$ then $u(x; A) \geq u(x; \{g < \delta\}) \geq 1 - \frac{g(x)}{\delta}$, which implies that $u(x; A) > 0$ in some neighborhood of $\{0\}$: this leads to a contradiction.

4.2. A two-dimensional non-linear Gaussian system. Let us provide a more computational example for the application of the methods developed in this work. Let $\mathcal{X} = \mathbb{R}^2$ be endowed with the standard topology, and consider a dtMP with dynamics given by the following system of non-linear difference equations:

$$(4.3) \quad \begin{cases} X_{1,n+1} &= 0.5X_{2,n}(3X_{1,n}^2 + 2X_{2,n}^2 - 0.5) + 0.6\eta_n \sqrt{X_{1,n}^2 + X_{2,n}^2}, \\ X_{2,n+1} &= 0.9X_{2,n}(2X_{1,n}^2 + 4X_{1,n}X_{2,n} + 3X_{2,n}^2 - 0.5) + 0.6\zeta_n \sqrt{X_{1,n}^2 + X_{2,n}^2}, \end{cases}$$

where $(X_{1,0}, X_{2,0}) = (x_1, x_0) = x$. Here $(\eta_n)_{n \geq 0}$ and $(\zeta_k)_{k \geq 0}$ are independent sequences of iid standard normal random variables. The process is weakly continuous with the origin $\{0\}$ as the absorbing set and its kernel can be expressed explicitly as in Section

4.1. We are interested to solve the infinite-horizon invariance problem over the compact set $A = [-0.6, 0.6] \times [-0.6, 0.6]$. Since $\text{l.a.s.}(A) = \{0\}$, the set A is non-simple, thus Theorem 3 cannot be applied as discussed in Section 4.2. It is thus necessary to find a δ -locally excessive function on A . Let us start by discussing the behavior of the process X on the phase plane. For x far from the origin, the non-linear terms (appearing in brackets in (4.3)) play a more important role than the linear ones, whereas for x close to the origin the situation is reversed. We then expect that a function measuring the distance from the origin may be locally excessive. For this reason, we consider $g(x) = \|x\|^2$, which leads to the following form for $\mathcal{P}g$:

$$(4.4) \quad \begin{aligned} \mathcal{P}g(x_1, x_2) = & \frac{1}{200}(144x_1^2 + 197x_2^2 - 474x_1^2x_2^2 + 1098x_1^4x_2^2 - 648x_1x_2^3) \\ & + \frac{1}{200}(2592x_1^3x_2^3 - 586x_2^4 + 5136x_1^2x_2^4 + 3888x_1x_2^5 + 1658x_2^6). \end{aligned}$$

It holds that $\{g < 0.25\} \subseteq \mathcal{E}_g$, hence g is δ -locally excessive on A , with $\delta = 0.25$. Figure 1 shows sets $\{g < 0.25\}$, \mathcal{E}_g , and A . Set A intersects \mathcal{E}_g^c , which can be interpreted as follows: starting from a state $x \in A \cap \mathcal{E}_g$, process X exhibits contractive dynamics, whereas for $x \in A \cap \mathcal{E}_g^c$ the difference $\mathcal{P}g(x) - g(x)$ is positive and gets larger as $\|x\|$ grows, hence trajectories initialized in $x \in A \cap \mathcal{E}_g^c$ expand away from the origin. Based on this consideration we expect clear differences between the values of function $u(x; A)$ for $x \in A \cap \mathcal{E}_g$ and $x \in A \cap \mathcal{E}_g^c$.





FIGURE 1. Infinite-horizon invariance problem over set A . The boundaries of sets \mathcal{E}_g (dark blue curves), A (brown square) and $\{g < 0.25\}$ (cyan circle).


We apply the decomposition technique in Theorem 2, where by selecting an $\varepsilon = 0.02$ we obtain that $0 \leq w(x; A, \{g < 5 \cdot 10^{-3}\}) - u(x; A) \leq 0.02$. To simplify the calculations, we consider the function $w(x; A, B)$ for $B = (-0.05, 0.05) \times (-0.05, 0.05) \subset \{g < 5 \cdot 10^{-3}\}$ and as a result we have $0 \leq w(x; A, B) - u(x; A) \leq 0.02$. To compute the values of function w we use the bounds provided in Proposition 2: in this case $m(A \setminus B) = 1$ and $\alpha(A \setminus B) \approx 0.957$, so by considering $n = 50$ iterations we obtain $0 \leq w(x; A, B) - w_n(x; A, B) \leq 0.105$.

Thus far, the methods developed in this paper (in particular Theorem 2), have allowed us to reduce the infinite-horizon invariance problem over a non-simple set to a finite-horizon reach-avoid problem. Let us now mention how the value function corresponding to the latter problem can be computed. The calculation of the value function w_n is performed with a target error 0.1, which is achieved by employing a standard uniform discretization algorithm [AKLP10] – thus the resulting overall error equals to 0.207. Based on the time horizon of the problem, and due to the degenerate nature of

the kernel in the neighborhood of the origin, and the fine size selected by the partitioning procedure to achieve the small required precision, the computation took 24 hours on Intel Core i5, 2.4 GHz with 4Gb RAM. This computational time can be further reduced by leveraging more involved numerical procedures [SA11], which however are outside of the scope of this study.



(a) Local excessivity of g on the set A



(b) Invariance value function

FIGURE 2. Results for the infinite-horizon invariance problem on the set A . Graphs of functions $\mathcal{P}g - g$ (a) and u (b).

The first goal of this case study was to show that a infinite-horizon problems can be solved efficiently, with strict bounds on the error, even in the case of nonlinear dynamics and kernels which admit non-trivial absorbing sets. The use of the decomposition technique has also allowed us to avoid computations over the neighborhood of the absorbing set $(0,0)$ where the kernel P degenerates. In particular, it is important for numerical methods based on the discretization of the state space, since their error bounds depends on Lipschitz constants of densities. Moreover, with this approach the error of computation can be made as small as needed by varying the error ε related to the decomposition, the number n of iterations for the reach-avoid problem, and the grid size for the discretization.

As already mentioned, the choice of the set A with regards to the excessive region plays an important role. On Figure 2(a) one can observe large positive values of $\mathcal{P}g(x) - g(x)$ for x close to points $(-0.6, 0.6)$ or $(0.6, 0.6)$. We expect a diverging behavior of X when starting in that region. This fact is clearly shown on Figure 2(b) where the invariance value function u takes the smallest values exactly in that region.

5. CONCLUSIONS

This work has provided a general framework for the study of formal algorithms for PCTL verification of discrete-time Markov processes over general state spaces. The main focus of the article has been placed on the verification of infinite-horizon PCTL specifications, both in terms of characterization of the given PCTL formula and in terms of precise numerical computation of the corresponding value function. It has been shown that structural properties of the stochastic kernel, namely the possible presence of absorbing subsets of given sets, are crucial for problems over the infinite horizon. In particular, the solution of the invariance is either trivial (on simple sets) or extremely complicated (on non-simple sets). This has led to criteria to distinguish such instances and to techniques to tackle the latter case – these techniques have been illustrated by two case studies.

The outcome of this work is that infinite-horizon problems cannot in general be solved exactly or algorithmically. However, precise reduction of these problems to finite-horizon analogues allows tapping on techniques for the latter, thus inheriting their scalability. This leads to an emphasis on the verification of the simplicity of a given set and on the development of procedures to find δ -locally excessive functions.

These questions represent compelling goals to the authors and are to be further pursued in future work, along with the application of the developed methods to other classes of specifications (beyond PCTL). Furthermore, extensions to continuous-time and control-dependent models are also deemed research worthy.

REFERENCES

- [AKLP10] A. Abate, J.-P. Katoen, J. Lygeros, and M. Prandini, *Approximate model checking of stochastic hybrid systems*, European Journal of Control **16** (2010), 624–641.
- [AKM11] A. Abate, J.-P. Katoen, and A. Mereacre, *Quantitative automata model checking of autonomous stochastic hybrid systems*, Proceedings of the 14th ACM international conference on Hybrid Systems: Computation and Control, 2011, pp. 83–92.
- [APLS08] A. Abate, M. Prandini, J. Lygeros, and S. Sastry, *Probabilistic reachability and safety for controlled discrete time stochastic hybrid systems*, Automatica **44** (2008), no. 11, 2724–2734.
- [BK08] C. Baier and J.-P. Katoen, *Principles of model checking*, The MIT Press, 2008.
- [BS78] D.P. Bertsekas and S.E. Shreve, *Stochastic optimal control: The discrete time case*, vol. 139, Academic Press, 1978.
- [DGJP04] J. Desharnais, V. Gupta, R. Jagadeesan, and P. Panangaden, *Metrics for labelled Markov processes*, Theoretical Computer Science **318** (2004), no. 3, 323–354.
- [Dur04] R. Durrett, *Probability: Theory and examples - third edition*, Duxbury Press, 2004.
- [FKNP11] V. Forejt, M. Kwiatkowska, G. Norman, and D. Parker, *Automated verification techniques for probabilistic systems*, Formal Methods for Eternal Networked Software Systems (2011), 53–113.
- [HKNP06] A. Hinton, M. Kwiatkowska, G. Norman, and D. Parker, *PRISM: A tool for automatic verification of probabilistic systems*, Tools and Algorithms for the Construction and Analysis of Systems (H. Hermanns and J. Palsberg, eds.), Lecture Notes in Computer Science, vol. 3920, Springer Verlag, 2006, pp. 441–444.
- [HLL96] O. Hernández-Lerma and J. B. Lasserre, *Discrete-time Markov control processes*, Applications of Mathematics (New York), vol. 30, Springer Verlag, New York, 1996. MR 1363487 (96k:93001)
- [Hut05] M. Huth, *On finite-state approximants for probabilistic computation tree logic*, Theoretical Computer Science **346** (2005), no. 1, 113–134. MR 2182232 (2006g:68171)
- [Kal02] O. Kallenberg, *Foundations of modern probability*, second ed., Probability and its Applications (New York), Springer Verlag, New York, 2002. MR 1876169 (2002m:60002)
- [KKZ05] J.-P. Katoen, M. Khattri, and I. S. Zapreev, *A Markov reward model checker*, Quantitative Evaluation of Systems (QEST), 2005, pp. 243–244.
- [Kus67] H.J. Kushner, *Stochastic stability and control*, vol. 33, Academic Press, 1967.
- [Mey08] S.P. Meyn, *Control techniques for complex networks*, Cambridge University Press, 2008.
- [MT93] S.P. Meyn and R.L. Tweedie, *Markov chains and stochastic stability*, Springer Verlag, 1993.
- [Pap03] G.J. Pappas, *Bisimilar linear systems*, Automatica **39** (2003), no. 12, 2035–2047. MR 2143530 (2005k:93022)
- [PS06] G. Peškir and A. Shiryaev, *Optimal stopping and free-boundary problems*, Birkhauser, 2006.

- [RCSL10] F. Ramponi, D. Chatterjee, S. Summers, and J. Lygeros, *On the connections between PCTL and dynamic programming*, Proceedings of the 13th ACM international conference on Hybrid Systems: Computation and Control, 2010, pp. 253–262.
- [Rev84] D. Revuz, *Markov chains*, second ed., North-Holland Publishing, Amsterdam, 1984. MR 758799 (86a:60097)
- [Rud76] W. Rudin, *Principles of mathematical analysis*, vol. 275, McGraw-Hill New York, 1976.
- [Rud87] ———, *Real and complex analysis*, third ed., McGraw-Hill Book Co., New York, 1987. MR 924157 (88k:00002)
- [SA11] S.E.Z. Soudjani and A. Abate, *Adaptive gridding for abstraction and verification of stochastic hybrid systems*, Quantitative Evaluation of Systems (QEST) (Aachen, DE), 2011, pp. 59–68.
- [SA12] ———, *Probabilistic invariance of mixed deterministic-stochastic dynamical systems*, Proceedings of the 15th ACM international conference on Hybrid Systems: Computation and Control (Beijing, PRC), 2012, pp. 207–216.
- [SL10] S. Summers and J. Lygeros, *Verification of discrete time stochastic hybrid systems: A stochastic reach-avoid decision problem*, Automatica **46** (2010), no. 12, 1951–1961.
- [SRG08] A. Shiryaev, B. Rozovskii, and G. Grimmett, *Optimal stopping rules*, Springer Verlag Berlin Heidelberg, 2008.
- [TA11] I. Tkachev and A. Abate, *On infinite-horizon probabilistic properties and stochastic bisimulation functions*, Proceedings of the 50th IEEE Conference on Decision and Control and European Control Conference (Orlando, FL), December 2011, pp. 526–531.
- [TA12] ———, *Regularization of Bellman Equations for Infinite-Horizon Probabilistic Properties*, Proceedings of the 15th International Conference on Hybrid Systems: Computation and Control (Beijing, PRC), April 2012, pp. 227–236.
- [Tab09] P. Tabuada, *Verification and control of hybrid systems: A symbolic approach*, Springer Verlag, New York, 2009. MR 2521445 (2010h:93001)

6. APPENDIX

Theorem 1 requires the compactness of the set A and the weak continuity of the kernel P , however some of the relations between statements in this theorem are true in the general case as it is shown in Figure 6. First of all, for the pair $1) \Leftrightarrow 2)$ it is




FIGURE 3. Generalization of the relations between statements of Theorem 1.

clear that $2)$ is a stronger statement in general. Moreover, from the proof of Theorem 1 it clearly follows that $1) \Rightarrow 2)$ without any assumptions on A and P . For $3) \Leftrightarrow 4)$ the following holds:

Proposition 5. ⁵ Equation (3.5) admits a unique solution if and only if $u(\cdot; A) \equiv 0$.

Proof. Equation (3.5) is linear, so if its solution is unique it is the trivial zero solution. Since $u(\cdot; A)$ is one of solutions, $u(\cdot; A) \equiv 0$.

Conversely, let us suppose that $u(\cdot; A) \equiv 0$ and let $f \in \mathbb{B}$ be any other solution of (3.5), so that $\|f\| > 0$. Clearly, the function $\tilde{f} := \frac{f}{\|f\|}$ is also a solution of this equation and $\tilde{f} \leq 1$. As a result, it holds that $\tilde{f} \leq u = 0$ (see Remark 1) so that $\tilde{f} \leq 0$. On the other hand, $-\tilde{f}$ is also a solution of (3.5) due to the linearity of the equation and $-\tilde{f} \leq u \equiv 0$ which leads to $\tilde{f} = 0$. However, we have $\|\tilde{f}\| = 1$ by definition, hence we come to a contradiction. \square

⁵This proposition generalizes a result from [RCSL10, Proposition 9], where the trivial invariance was shown to be sufficient for the uniqueness over a smaller class of functions.

Now we only left to discuss relations between 2), 3) and 5). From contraction mapping theorem it follows that $2) \Rightarrow 3)$. Moreover, if $u(\cdot; A) \equiv 0$ then A is simple since $\text{l.a.s.}(A) = \{u(\cdot; A) = 1\}$ is empty in this case. As a result, all the relations in Figure 6 are true. Let us provide examples that other relation does not hold when either A is not compact or P is not weakly continuous. We first show that the weak continuity is not sufficient.

1. Let us show that $3)+5) \not\Rightarrow 2)$. Consider an example from Section 4.1 given by the equation (4.2) with $\mu = 0$ and $h(0, \sigma) \geq 0$. Let us choose the set $A = [-1, 1]$ so as it has been proved, $u(x; A) = 1_{\{0\}}(x)$. Let us put $\tilde{A} = A \setminus \{0\}$, i.e. it is not compact. By induction it can be proved that $u_n(x; A) - u_n(x; \tilde{A}) = 1_{\{0\}}(x)$ for all $n \geq 0$ since it holds for $n = 0$ and

$$\begin{aligned} u_{n+1}(x; A) - u_{n+1}(x; \tilde{A}) &= 1_A(x) \int_{\mathcal{X}} u_n(y; A) P(x, dy) - 1_{\tilde{A}}(x) \int_{\mathcal{X}} u_n(y; \tilde{A}) P(x, dy) \\ &= 1_{\{0\}}(x) + 1_{\tilde{A}}(x) \int_{\{0\}} P(x, dy) = 1_{\{0\}}(x). \end{aligned}$$

Note, however that if $f \in \mathbb{B}$ is continuous on $A = [-1, 1]$ so is $\mathcal{S}_A f$ due to the structure of the kernel \mathcal{P} . As a result, functions $u_n(\cdot; A)$ are continuous on A and since $u_n(0; A) = 1$ for all $n \geq 0$ it holds that $\|u_n(\cdot; \tilde{A})\| = \|u_n(\cdot; A) - 1_{\{0\}}(\cdot)\| = 1$. Although the set \tilde{A} is simple, its invariance value function $u(\cdot; \tilde{A}) \equiv 0$ and the uniqueness for the solution of (3.5) holds, we have $m(A) = \infty$ which proves that $3)+5) \not\Rightarrow 2)$ in general.

2. Let us show that $5) \not\Rightarrow 3)$. Following the same lines as above, we consider (4.2) with $\mu = 0$ and $h(0, \sigma) < 0$. As it has been discussed in Section 4.1 the function $u(x; A)$ for $A = [-1, 1]$ is positive in the neighborhood of 0. Still, it holds that $u(x; \tilde{A}) = u(x; A) - 1_{\{0\}}(x)$ for $\tilde{A} = A \setminus \{0\}$ so the invariance value function for the simple non-compact set \tilde{A} is non-trivial and hence the solution of (3.5) is not unique.

Finally, we can show that if A is compact but the weak continuity assumption on P is relaxed then $3)+5) \not\Rightarrow 2)$ and $5) \not\Rightarrow 3)$. To do this, one should make similar considerations as in **1.** and **2.** for the same kernels, just redefining $P(0, \{2\}) = 1$.